Automatic Registration of Mammograms Based on Linear Structures

Robert Marti[1], Reyer Zwiggelaar[1] and Caroline Rubin[2]

[1] Division of Computer Science, University of Portsmouth,Porstmouth, UK
{robert.marti,reyer.zwiggelaar}@port.ac.uk
[2] Breast Screening Unit, Royal South Hants Hospital, Southampton, UK

Abstract. A novel method to obtain correspondence between landmarks when comparing pairs of mammographic images from the same patient is presented. Our approach is based on automatically established correspondence between linear structures (i.e. ducts and vessels) which appear in mammograms using robust features such as orientation, width and curvature extracted from those structures. In addition, a novel multiscale feature matching approach is presented which results in a reliable correspondence between extracted features.

1 Introduction

Detection of abnormal structures or architectural distortions in mammographic images can be performed by comparing different images of the same patient, either the same breast taken at different times (temporal comparison) or using mammographic images of the left and right breast (contralateral comparison). This comparison is not straightforward due to additional dissimilarities between images which are related to patient movement, sensor noise, different radiation exposure and variation of breast compression specially as 2D mammographic images are projections of 3D mammographic structures. Therefore, in order to efficiently compare two mammograms and avoid non target dissimilarities, an initial alignment (also referred to as registration) must be carried out. Methods that are able to recover local deformation (e.g. [1]) rely on corresponding landmarks between images, which turns out to be the most difficult task and plays an important role in registration accuracy. Manually landmark generation is a tedious and time consuming task when the number of control points is large and, moreover, introduces variability. Automatic landmarking methods are, therefore, more suitable but also difficult to develop. Automatically extracted mammographic landmarks include breast boundary [4], pectoral muscle [4], salient regions [5] and crossings of horizontal and vertical structures [7].

This paper presents a novel method to establish image correspondence in mammographic images based on matching their major linear structures (ducts and vessels). Establishing correspondence involves various steps: 1. to identify linear structures in both mammograms (section 2), 2. to extract reliable information from those structures (section 3), 3. to obtain correspondence between the structures (section 4) and 4. registration using a point based method [1].

2   Detection of Linear Structures

We use a non-linear line operator [3] to detect linear structures in both mammograms. At a given scale, the line operator provides for every pixel a strength and orientation of the linear structure. Line strength is obtained by comparing a straight line of pixels with its neighbourhood and obtaining the maximum value from a number of orientations. Direction of the line strength is determined by the angle of the line which gives the highest line strength. Scale information is obtained from the maximum line strength of the detector at different scales.

2.1   Scale Information

In order to obtain more reliable scale information we investigate the feasibility of normalising the individual line strength images before the scale information is extracted. For the evaluation we have used a dense mammographic background to which synthetic linear structures of different scales have been added. Table 1 compares the scale information obtained with (a) no normalisation, (b) histogram stretching and (c) histogram stretching with the top 1% of the grey-levels mapped to the maximum grey-level value. Scale information using non-processed line strength images results in a biased scale estimation towards the higher scales and detected lines often overestimate the scale information. On the other hand, maximum mapping in combination with stretching provides more reliable scale information with no bias as shown in table 1c. It is our understanding that without normalisation the resulting line strength at the lower scales is affected (and relative to the higher scales suppressed) by relative noise levels.

Table 1. Scale detection using (a) no normalisation (45% detected), (b) histogram stretching (51% detected) and (c) histogram stretching with the top 1% of the grey-levels mapped to the maximum grey-level value (87% detected)

(a)		True scale			(b)		True scale			(c)		True scale		
		1	2	4			1	2	4			1	2	4
(a)	1	0.04	0.00	0.00	(b)	1	0.22	0.00	0.00	(c)	1	0.33	0.09	0.02
Detected	2	0.28	0.12	0.00	Detected	2	0.11	0.00	0.00	Detected	2	0.00	0.27	0.00
scale	4	0.01	0.27	0.29	scale	4	0.01	0.38	0.29	scale	4	0.00	0.02	0.27

The fact that scale information is improved when using the proposed normalisation is clear, but it is also important to analyse how this normalisation affects detection results, that is, the resulting line strength image. Receiver operating characteristic curves obtained by comparing true structures with the ones detected showed that no processing of strength images gives slightly better detection results, but this is only a small percentage and not comparable to the improved scale detection as indicated in table 1.

2.2   Line Processing

Once strength, direction and scale information have been obtained we perform different operations to facilitate the feature extraction process. First, we set a conservative threshold on the line strength image in order to remove background noise. Then non-maximum suppression is applied which removes pixels with low intensity values compared to their neighbours along the normal of the linear structure. Scale information extracted from the line operator (see section 2.1) is used here to determine the position of candidate pixels to be suppressed. In addition, short lines which do not provide reliable information are removed taking the Euclidean distance between the centre pixel and its neighbours into account. Finally a thinning operation will obtain the backbone of the most representative linear structures in the mammograms.

3   Feature Extraction

Feature extraction is needed in order to obtain descriptors of the structures to be used in the matching process. Corresponding linear structures in two mammograms can present large differences related to line strength and line continuity (due to different imaging conditions) but width and orientation of the line and local curvature and branching points are more likely to be preserved and often are features used by radiologists when comparing mammograms. Therefore, features which take line length, end points and line strength into account turn out to be unreliable features to tackle the correspondence problem. In this paper we use local features such as curvature, width and orientation. The basic idea of our method is to extract characteristic points of linear structures determined by the maximum curvature along the structure. Position, orientation (ϕ) and width (w) are then extracted for those points and used in the matching process.

Curvature measures are extracted for each pixel along the linear structures. Maximum curvature points are likely to be characteristic for a linear structure in terms of local curvature and branching points. Before computing curvature we need to extract the orientation of the structures which is obtained directly from the thinned linear structures. Although orientation information could be retrieved from the line detector results, experiments have shown (results not included) that the approach adopted here gives more accurate orientation measures. Curvature values at each pixel are obtained with a similar approach as used in [2]. Curvature (or directional change) between two pixels p and q is defined by the scalar product of their normal vectors. Hence, the curvature measure of a given pixel p is obtained by computing the scalar product between p and its neighbouring pixels.

$$C_p = \frac{1}{N} \sum_{i=1}^{N} \exp(-d_{ip}^2)(1 - \cos(\phi_p - \phi_i)) \tag{1}$$

where ϕ_i is the angle of the normal at each pixel i. As we will be extracting curvature from binary thinned images, we assume unit vectors. N is the number

of points in a local neighbourhood and d_{ip} is the Euclidean distance between points i and p. The distance factor is used here to weight the curvature of each point i, in order to incorporate a bias to points closer to p.

Width information is extracted after non-maximum suppression of the strength images (section 2.2). The improved scale information from the line detection step (section 2.1) is indirectly used to extract width information as scale information is used to perform the non-maximum suppression. Width of a linear structure at a point is given by the number of pixels along the normal of the structure.

4   Matching

The matching process needs to consider the following assumptions:

– Non-rigid motion: linear structures in mammograms suffer local distortions, therefore they may move independently and no geometrical relationship is established between neighbouring structures.
– Multiple matches: a linear structure in one mammogram can match more than one structure in the other mammogram, and vice versa.
– Non-bijectivity: a linear structure in one mammogram may not have a corresponding linear structure in the other, and vice versa.
– Localisation: After global breast misalignment is removed, matched linear structures lie in approximately the same area in both mammograms. We will refer to this area as the localisation area M.

We adopt here a similar but more general approach than the one used in [7]. We denote the set of feature points from both mammograms, as $\{a_i | 1 \leq i \leq N_i\}$ and $\{b_j | 1 \leq j \leq N_j\}$, where N_i and N_j are the number of feature points used, which may not be the same. Subsequently, we build a distance matrix (DM) in which each position $DM(i, j)$ describes the normalised distance between features of points a_i and b_j. Hence, a low value means good matching between points. The computation of the feature distance will be discussed later. Once all distances have been entered, the minimum value of each row is detected (and remaining positions in each row deleted) in order to have a unique match in each row. Then, the minimum value of each column is extracted ending up with a set of potentially matched points. It must be mentioned that the approach only works when there are distinct minima in each row and column and that points that do not conform to this are removed from the data-set. Reliable matches are those with a distance smaller than a particular threshold. The use of the distance matrix structure fullfills the first three assumptions: independent motion (matched points a_i, b_j do not imply matching a_{i+1}, b_{j+1}); a point a_i may have multiple matched points b_j; and a point in either mammogram may remain unmatched.

As mentioned previously the distance matrix contains the normalised distance between features of points a_i and b_j. Satisfying the last assumption, localisation, position $DM(i, j)$ will only have a finite value if points a_i and b_j are in the same localisation area in both mammograms. This assumption can only be stated if mammograms are globally aligned, that is, global deformation (i.e. rotation,

translation and scale) is removed. Therefore, we initially register mammograms maximising a mutual information measure using an affine transformation [6]. Registration will provide the transformation parameters (α) needed to compare feature points coordinates and establish the localisation area (M) in both mammograms. The normalised distance is determined by three components. The first distance is the Euclidean distance (D_E) between point coordinates. Coordinates of one of the points are transformed (T_α) using the parameters obtained from the registration mentioned earlier. The second distance is the orientation difference between two points (D_ϕ). Finally, the last distance is the width difference between two points (D_w) normalised using the maximum width of all the linear structures (W). Using equal weighting the normalised distance is given by

$$
\begin{aligned}
DM(i,j) &= D_\phi + D_w + D_E \\
&= \frac{1 - \cos(\phi_i - \phi_j)}{2} + \frac{|w_i - w_j|}{W} + \frac{|T_\alpha(a_i) - b_j|}{M}.
\end{aligned}
\tag{2}
$$

4.1   Multi-level Matching

The described matching process may obtain matching points which lie localised in small areas and not spread over the whole image, as it would be necessary for image registration. In addition, the localisation area defined using the global registration may not be accurate enough and a local registration should be used. A novel multi-level registration approach is used here to tackle those problems. At the first level, the full images are registered obtaining the transformation parameters. Subsequently, we move to the second level dividing each mammogram in six rectangular sub-images and again register each sub-image. Transformation parameters are carried through each level, assuming that each sub-image at lower levels would suffer a different transformation but it would be related to the deformation on the higher level. Assuming this, we speedup the optimisation process as well as avoiding local minimum situated away from the optimum solution. Once the last level is reached, transformation parameters in each sub-image on that level establish the correspondence of localisation areas for structures within each sub-image. In addition, extracting the local best matches in each sub-image assures that a minimum number of matches will be present, having a more homogeneous point distribution over the whole mammogram.

5   Results

In this section we present initial results using the described approach applied to temporal and contralateral mammographic comparison. Figure 1 shows two mammograms of the same patient taken three years apart where matches between the linear structures are indicated by numbers. As mentioned earlier, matched points can be used as control points in mammographic registration, using a point based method such as thin plate splines [1]. Figure 1c shows the

subtracted image (where darker areas mean larger misalignment) obtained after automatic registration using the proposed method. Although mis-registration can be observed near the breast outline using our method, registration of internal breast regions is comparable to manual placement of control points. This statement is corroborated in figure 2 which shows temporal and contralateral registration results using our method compared to manually placed control points. Graphs are obtained measuring normalised mutual information between reference and registered images. A high value denotes high similarity between images, therefore good registration. Graph results show that automatic registration performs equally or slightly better in most cases, although some poor results are also obtained. These are due to specific breast characteristics such as the lack of major linear structures (mammograms 4 and 5 in figure 2a) or large image deformation (mammograms 15 and 16 in figure 2b).

Fig. 1. Correspondence in temporal mammograms (from left to right): reference image, warped image, difference between registered image and reference
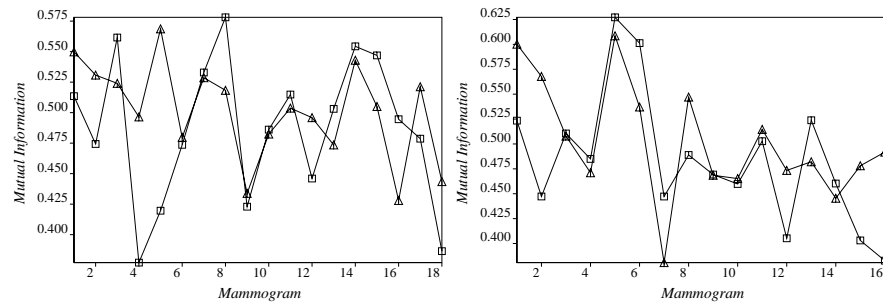
Fig. 2. Registration results for temporal (left) and contralateral (right) experiments where automatic registration (□) is compared to manual registration (△)

6   Conclusions

The work presented here describes a novel approach to solve the problem of extraction of reliable features in mammographic images and establishes correspondence between them in pairs of mammograms. We have shown that features extracted from linear structures can provide an automatic approach to the generation of control points for image registration. Features based on scale, orientation and position have been used. Initial results look promising, but further work will be needed to establish the full benefit of our approach.

The proposed method will be tested on a larger mammographic dataset and compared to a radiologist gold standard. In addition, other features could be incorporated such as breast boundary and the position of the nipple.

References

1. F. Bookstein. Principal warps: thin-plate splines and the decomposition of deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(6):567–585, 1989.
2. J. Deschênes and D. Ziou. Detection of line junctions and line terminations using curvilinear features. *Pattern Recognition Letters*, 21:637–649, 2000.
3. R. Dixon and C. Taylor. Automated asbestos fibre counting. *Inst. Phys. Conf. Ser*, 44:178–185, 1979.
4. N. Karssemeijer and G.M. te Brake. Combining single view features and asymmetry for detection of mass lesions. 4^{th} *International Workshop on Digital Mammography*, Nijmegen, The Netherlands, pp. 95-102, 1998.
5. S. Kok-Wiles, J.M. Brady and R.P. Highnam. Comparing mammogram pairs in the detection of lesions. 4^{th} *International Workshop on Digital Mammography*, Nijmegen, The Netherlands, pp. 103-110, 1998.
6. R. Marti, R. Zwiggelaar, and C. Rubin. Comparing image correspondence in mammograms. 5^{th} *International Workshop on Digital Mammography*, 2000. In press.
7. N. Vujovic and D. Brzakovic. Establishing the correspondence between control point in pairs of mammographic images. *IEEE Transactions on Image Processing*, 6(10):1388–1399, 1997.